

AMD Opteron™ Processors

A Better High-End Embedded Solution

David Jessel
AMD Boston Design Center

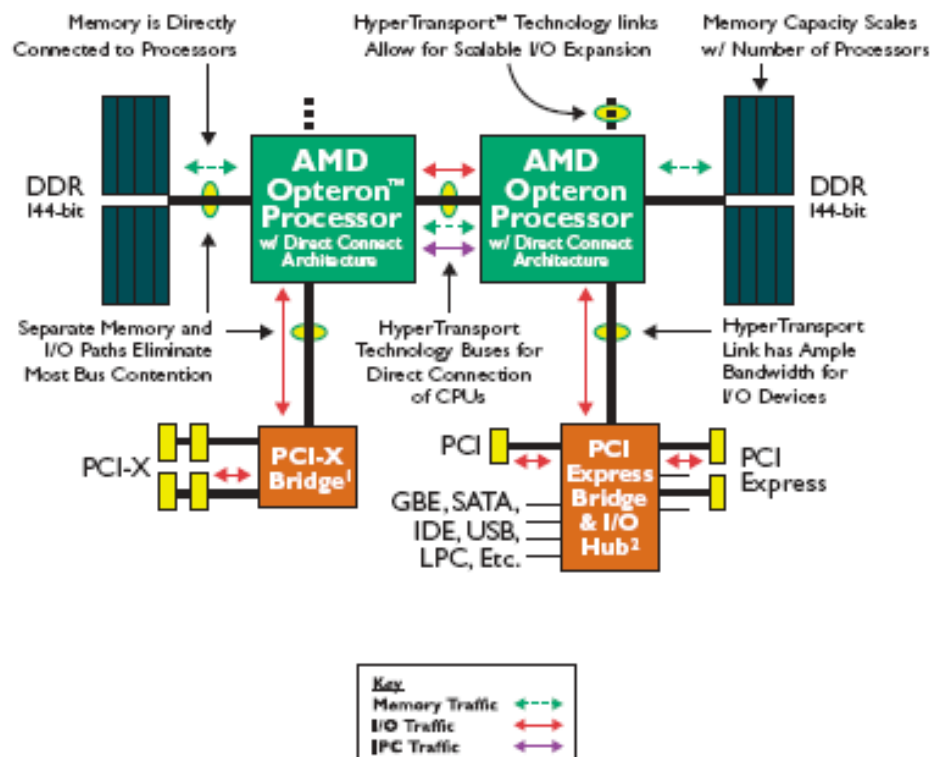
© 2005 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD Arrow logo, AMD Opteron, AMD Geode, AMD Alchemy, AMD Athlon64 and combinations thereof are trademarks of Advanced Micro Devices, Inc. HyperTransport is a licensed trademark of the HyperTransport Technology Consortium. PCI-X is a registered trademark of PCI-SIG. Other product and company names in this paper are for identification purposes only and may be trademarks of their respective companies.

AMD has a two-pronged approach to the embedded market: addressing the high-end with AMD Opteron™ processors; addressing the low-end with AMD Geode™ and AMD Alchemy™ solutions. The following sections explain how the balanced and scalable architecture of AMD Opteron processors satisfies the most demanding performance needs of embedded applications like digital media, storage, and networking. Embedded application performance requirements can vary from CPU, memory, and I/O throughput (to analyze MRI images and large scientific datasets or stream HDTV video) to secure scientific and enterprise storage data to delivering networking information (voice, data, video, wired or wireless). These applications usually require longevity of supply and may require dense enclosures where lower power processing is needed. These requirements are met with AMD Opteron processors.

AMD Opteron™ processors are well suited for the embedded market with Direct Connect Architecture, low power alternatives, and supply longevity for certain models. AMD's Direct Connect Architecture directly connects the processor to memory, processor to I/O, and processor to other processors through coherent HyperTransport™ technology links. Direct Connect Architecture provides leadership performance for the embedded market by improving bandwidth, removing bottleneck contentions, and reducing latency.

The following figure is a view of the AMD Opteron™ processor Direct Connect Architecture in a dual processor configuration:



Direct Connect Architecture -- Memory

With Direct Connect Architecture, each AMD Opteron™ processor has an integrated memory controller directly connected to DDR memory. This dramatically decreases latency and improves memory bandwidth.

Multiprocessor AMD Opteron-based systems have leadership memory bandwidth and can support NUMA operating systems in addition to Symmetrical Multiprocessing (SMP).

Direct Connect Architecture implemented in the AMD Opteron processor is superior to legacy Northbridge architecture where the processor communicates through a slow front side bus to a Memory Hub chip that then connects to DDR memory. This front side bus is constrained by memory, I/O, and processor-to-processor communications. The Memory Hub adds a component to the board which may limit the density of the board and the application that can be targeted. This memory hub also adds latency to memory. Northbridge architecture is inferior in many ways and application performance benchmarks prove this true. Using a revolutionary approach like AMD's Direct Connect Architecture solves these performance problems.

HyperTransport™ Technology for High-bandwidth, Low-latency I/O

AMD Opteron™ processors have integrated HyperTransport™ technology on die. HyperTransport technology is a clock-forwarded, point-to-point, high-bandwidth/low-latency interconnect. There are three 16-bit HyperTransport technology links on the AMD Opteron, and one HyperTransport technology link on AMD Athlon64™ processors. Each 16-bit link runs at 1GHz dual data rate (2Gbps). Each link is two bytes, which means each link is running at 4GBytes/s in each direction (32 Gbps in each direction) or 8 GBytes/s total. HyperTransport technology is a very efficient and sustains 80% of the peak performance (25Gbps or 3.2GBytes/s).

AMD's chipsets connect via HyperTransport technology to industry-standard protocols such as PCI-X, PCI-X 2.0, and PCI-Express. A Southbridge is used to connect to flash and legacy protocols. There are single-chip solutions that are especially suitable for embedded applications.

The HyperTransport Technology Consortium has over 50 member companies. These include leading chip, systems, storage, and networking companies like Apple, Broadcom, Cisco, HP, IBM, PMC Sierra, Sun, Transmeta, nVidia, and AMD.

Visit www.hypertransport.org/ for more information.

Direct Connect Architecture – Processor to Processor

Coherent HyperTransport technology provides glueless multi-processing, making it possible for system designers to simply interconnect AMD Opteron™ processors without any glue logic in between. AMD Opteron 800 series processors support up to an 8-way system; AMD Opteron 200 series processors supports up to a 2-way system; and the AMD Opteron 100 series processors supports a uni-processor system.

Low Power

AMD Opteron™ processors have power levels of 95W, 55W, and 30W to meet the performance and performance/watt needs of dense computing. These power levels deliver the highest performance at 95W where cooling is not an issue; offer 55W where density becomes more important; and 30W where the focus is on density first and CPU performance second. AMD is able to deliver more performance at these power levels than our competition because of AMD's focus on SOI process technology. The benefit of the AMD Opteron processor is that the same pin out can satisfy all three power levels. Functionality is not reduced to achieve this low-power; core voltage is lowered, and the parts are screened to meet the power specification.

Examples of 55W systems are standard servers and blades, where performance per watt is a key criteria. The systems are designed for high-density without sacrificing high-performance.

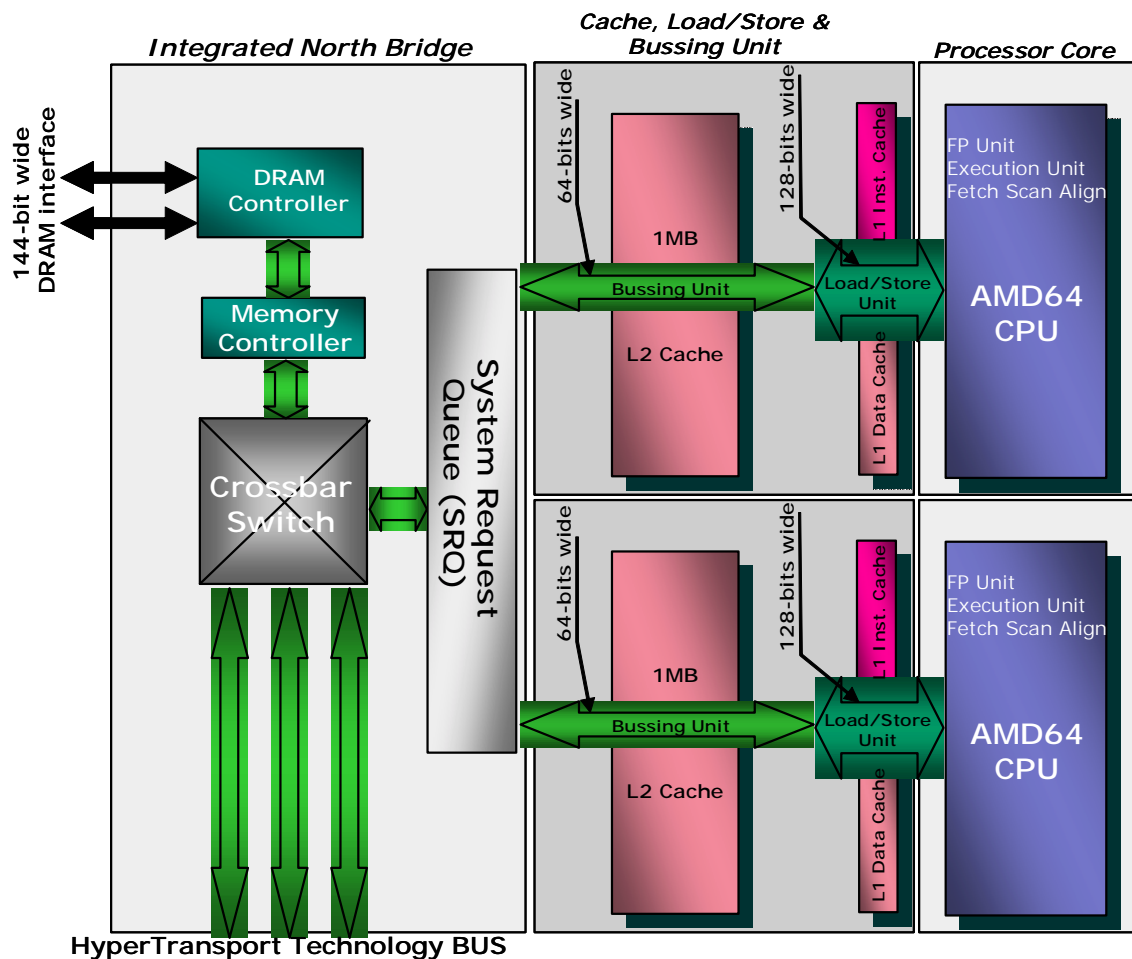
Examples of 30W systems are blades and switches where I/O performance per watt is paramount. The systems in networking storage and telecommunications can take advantage of the AMD Opteron processor's low-power and density and HyperTransport technology performance.

Dual Core

On August 31, 2004, AMD demonstrated the first x86-compatible dual core processor – the AMD Opteron™. Dual core processors provide the ability to significantly increase performance and throughput while maintaining the software compatibility that is expected from AMD-based processors. The benefit of dual core to the embedded market is obvious – higher performing processors in an even smaller footprint for improved density and performance.

In a dual core configuration, the two AMD Opteron processor CPU cores are connected via a System Request Queue (SRQ), designed from the beginning to be dual-core capable. Thus, the intelligence of the on-chip SRQ, crossbar switch, HyperTransport technology links, and memory controller efficiently optimizes the dual-core performance. Dual core AMD Opteron processors are expected to ship in volume by mid-2005.

The following is a block diagram of an AMD Opteron™ dual-core processor:



HyperTransport™ Technology Connector -- HTX

The HyperTransport™ technology HTX connector provides compelling high bandwidth with low latency solutions. HTX connectors deliver exceptional I/O functionality and flexibility by enabling direct connectivity between HyperTransport™ technology cards and AMD Opteron™ processors.

Consider an AMD Opteron processor-based 2P or 4P motherboard with a HyperTransport technology connector where an I/O expander card can be directly connected to one of the processor HyperTransport technology links. This expander board could be an accelerator card based on an FPGA or HyperTransport technology ASIC. The same expander board could also provide I/O functionality such as Infiniband, 10 Gigabit Ethernet with TOE, and fiber channel. This high-bandwidth low-latency HyperTransport technology link provides breakthrough advantage in the embedded market where I/O performance is critical.

Longevity

AMD is responding to customers' requests for long-term production to support embedded AMD Opteron™ processor-based solutions. AMD has selected a subset of the AMD Opteron processor product roadmap for five year steady-state production, followed by a two year end-of-life period. This longevity provides support for high-end embedded solutions. AMD Opteron processor Direct Connect Architecture and low-power models provide compelling solutions in storage, communications and networking, and imaging.

Ecosystem

AMD Opteron™ and AMD Athlon™64 processors implement the AMD64 architecture which supports x86 32-bit and the x86-64 standard developed by AMD. Thus, the x86 infrastructure of applications, compilers, development tools, operating systems, debuggers, etc. are all supported on AMD Opteron and AMD Athlon 64 processors. Examples of this infrastructure are based on embedded Linux, BSD, and Windows XP embedded operating systems.

Futures

There are designs underway with AMD Opteron™ and AMD Athlon™64 processors being used in 1U compute blades, networking, and storage blades. There are also AMD Opteron and AMD Athlon64 processor-based single board computing (SBC) designs. There are also efforts in Advanced TCA systems. These solutions will show that AMD Opteron and AMD Athlon64 processors are an excellent fit for embedded systems.

To learn more about the AMD Opteron Processor visit

www.amd.com/opteron

To read more AMD Opteron white papers visit

www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8796_8807,00.html

To learn more about AMD Alchemy™ Solutions and AMD Geode™ Solutions visit

www.amd.com/connectivitysolutions